

A FIELD GUIDE TO FORWARD-BACKWARD SPLITTING WITH A FASTA IMPLEMENTATION

TOM GOLDSTEIN, CHRISTOPH STUDER, RICHARD BARANIUK

ABSTRACT. Non-differentiable and constrained optimization play a key role in machine learning, signal and image processing, communications, and beyond. For high-dimensional minimization problems involving large datasets or many unknowns, the forward-backward splitting method (also known as the proximal gradient method) provides a simple, yet practical solver. Despite its apparent simplicity, the performance of the forward-backward splitting method is highly sensitive to implementation details.

This article provides an introductory review of forward-backward splitting with a special emphasis on practical implementation aspects. In particular, issues like stepsize selection, acceleration, stopping conditions, and initialization are considered. Numerical experiments are used to compare the effectiveness of different approaches.

Many variations of forward-backward splitting are implemented in a new solver called FASTA (short for Fast Adaptive Shrinkage/Thresholding Algorithm). FASTA provides a simple interface for applying forward-backward splitting to a broad range of problems appearing in sparse recovery, logistic regression, multiple measurement vector (MMV) problems, democratic representations, 1-bit matrix completion, total-variation (TV) denoising, phase retrieval, as well as non-negative matrix factorization.

CONTENTS

1. Introduction	2
1.1. Outline	2
1.2. A Word on Notation	3
2. Forward-Backward Splitting	3
2.1. Convergence of FBS	4
3. Applications of Forward-Backward Splitting	4
3.1. Simple Applications	5
3.2. More Complex Applications	7
3.3. Non-convex Problems	9
4. Bells and Whistles	11
4.1. Adaptive Stepsize Selection	11
4.2. Preconditioning	12
4.3. Acceleration	13
4.4. Backtracking Line Search	15
4.5. Continuation	16
4.6. Stopping Conditions	16
5. FASTA: A Handy Forward-Backward Solver	17
6. Numerical Experiments	18
6.1. Test Problem Details	18
6.2. Results and Discussion	20
7. Conclusion	21
Appendix A. Convergence Proof for Non-Monotone Line Search	22

Date: February 17, 2016.

1. INTRODUCTION

A large number of non-differentiable and constrained convex optimization problems have the following form:

$$(1) \quad \text{minimize} \quad h(x) = f(x) + g(x),$$

where $x \in \mathbb{R}^N$, f is convex and differentiable, and g is an arbitrary (i.e., not necessarily smooth) convex function. Problems of the form (1) arise frequently in the many areas including machine learning, dictionary learning [1], online learning for streaming data [2], sparse coding [3, 4], spectral clustering [5], sparse covariance estimation [6], approximate nearest neighbor search and vector quantization [7], as well as compressive sensing [8].

In many situations, the function g is neither differentiable nor even finite-valued, in which case the problem (1) cannot be minimized using simple gradient-descent methods. However, for a large class of functions g that arise in practice, one can efficiently compute the so-called *proximal* operator

$$(2) \quad \text{prox}_g(z, \tau) = \arg \min_x \tau g(x) + \frac{1}{2} \|x - z\|^2.$$

The proximal operator finds a point close to the minimizer of g without straying too far from a starting point z , and is often referred to as a *backward* (or *implicit*) gradient-descent step with stepsize τ . If the proximal operator (2) can be evaluated easily, then one can solve (1) efficiently using the *Forward-Backward Splitting* (FBS) method (also known as the proximal gradient method). Put simply, FBS can handle non-differentiable objectives and convex constraints while maintaining the simplicity of gradient-descent methods.

Due to the vast applications of FBS and its utility for sparse coding and regression, many variants of FBS have been developed to improve performance and ease of use. In its raw form, FBS requires the user to choose a number of convergence parameters that strongly effect both the performance and reliability of the algorithm. These include stepsizes, stopping condition parameters, acceleration schemes, stability conditions, and initialization. With the right modifications, FBS can be performed without substantial oversight from the user.

This article introduces and reviews FBS from a practical point of view. While this is not the first review article written on FBS, it differs from existing review articles in that it focuses on practical implementation issues. For an excellent theoretical review of FBS, we refer to [9].

1.1. Outline. In Section 2, we introduce the forward-backward splitting method and discuss its convergence behavior. Numerous example problems are discussed in Section 3, and for each we detail the formulation and solution via FBS. In Section 4, we discuss practical issues related to the implementation of FBS. The performance of variants of FBS on different test problems is explored in Section 6. Practical issues of FBS are discussed in Section 4 and are incorporated into a new reference implementation of FBS, called FASTA (short for Fast Adaptive Shrinkage/Thresholding Algorithm). FASTA provides a simple interface for applying forward-backward splitting to a broad range of optimization problems.

1.2. A Word on Notation. The ℓ_2 norm is denoted by $\|x\| = \sqrt{\sum_i |x_i|^2}$. We use the “single-bar” notation to define the ℓ_1 norm $|x| = \sum_i |x_i|$. The transpose of a matrix/vector x is denoted x^T . The inner-product of two vectors x and y is denoted $\langle x, y \rangle = x^T y$. The inner product of two matrices X and Y is given by $\langle X, Y \rangle = \text{trace}(X^T Y) = \sum_{i,j} X_{i,j} Y_{i,j}$. We use the symbol \Re to denote the real part of a (possibly) complex quantity. For example $\Re x$ denotes the real part of x . The vector x^* that minimizes some objective function h is denoted $x^* = \arg \min_x h(x)$. In contrast, the minimal *value* of f is denoted $\min_x f(x)$.

2. FORWARD-BACKWARD SPLITTING

Forward-Backward Splitting is a two-stage method that addresses each term in (1) separately. The FBS method is listed in Algorithm 1.

Algorithm 1 Forward-Backward Splitting

while not converged **do**
(3) $\hat{x}^{k+1} = x^k - \tau^k \nabla f(x^k)$
(4) $x^{k+1} = \text{prox}_g(\hat{x}^{k+1}, \tau^k) = \arg \min_x \tau^k g(x) + \frac{1}{2} \|x - \hat{x}^{k+1}\|^2$
end while

Let’s examine each step of the algorithm in detail. Line (3) performs a simple *forward* gradient descent step on f . This step begins at iterate x^k , and then moves in the direction of the (negative) gradient of f , which is the direction of steepest descent. The scalar τ^k is the *stepsize* which controls how far the iterate moves along the gradient direction during iteration k .

Equation (4) is called the proximal step, or *backward* gradient descent step. To understand this terminology, we examine the proximal operator (2). Any x^* that minimizes (2) must satisfy the optimality condition

$$(5) \quad 0 = \tau G + (x^* - z)$$

where $G \in \partial g(x^*)$ is some sub-gradient (generalized derivative) of g . Note that when g is differentiable we simply have $G = \nabla g(x^*)$. Equation (5) rearranges to

$$x^* = \text{prox}_g(z, \tau) = z - \tau G.$$

This shows that x^* is obtained from z by marching down the sub-gradient of g . For this reason, the proximal operator performs a gradient descent step. Because the sub-gradient G is evaluated at the final point x^* rather than the starting point z , this is called *backward* gradient descent.

Equation (5) is equivalent to the set inclusion $0 \in \tau \partial g(x^*) + (x^* - z)$, which rearranges to

$$z \in \tau \partial g(x^*) + x^* = (\tau \partial g + I)x^*.$$

For this reason, the proximal operator (2) is sometimes written

$$x^* = (\tau \partial g + I)^{-1} z = J_{\tau \partial g} z$$

where $J_{\tau \partial g} = (\tau \partial g + I)^{-1}$ is the *resolvent operator* of $\tau \partial g$. The resolvent is simply another way to express the proximal operator. In plain terms, the proximal/resolvent operator, when applied to z , performs a backward gradient descent step starting at z .

Algorithm 1 alternates between forward gradient descent on f , and backward gradient descent on g . The use of a backward step for g is advantageous in several ways. First, it may be difficult to choose a sub(gradient) of g in cases where the sub-gradient ∂g has a complex form or is not unique. In contrast, it can be shown that problem (4) always has a unique well-defined solution [10], and (as we will see later) it is often possible to solve this problem in simple closed form. Second, the backward step has an important effect on the convergence of FBS, which is discussed in the next section.

2.1. Convergence of FBS. The use of backward (as opposed to forward) descent for the second step of FBS is needed to guarantee convergence. To see this, let x^* denote a fixed point of the FBS iteration. Such a point satisfies

$$x^* = \text{prox}_g(x^* - \tau^k \nabla f(x^*), \tau^k) = x^* - \tau^k \nabla f(x^*) - \tau^k G(x^*)$$

for some $G(x^*) \in \partial g(x^*)$. This simplifies to

$$0 = \nabla f(x^*) + G(x^*),$$

which is the optimality condition for (1). This simple argument shows that a vector is a fixed-point of the FBS iteration if and only if it is optimal for (1). This equation has a simple interpretation: when FBS is applied to an optimal point x^* , the point is moved to a new location by the forward descent step, and the backward descent step puts it back where it started. Both the forward and backward step agree with one another because they both rely on the gradients evaluated at x^* .

This simple fixed point property is not enough to guarantee convergence. FBS is only convergent when the stepsize sequence $\{\tau^k\}$ satisfies certain stability bounds. An important property of FBS is that this stability condition does not depend on g , but rather on the curvature of f . In the case of a constant stepsize $\tau^k = \tau$, FBS is known to converge for

$$(6) \quad \tau < \frac{2}{L(\nabla f)}$$

where $L(\nabla f)$ is a Lipschitz constant of ∇f (i.e., $\|\nabla f(x) - \nabla f(y)\| < L\|x - y\|$ for all x, y). In many applications $f = \frac{1}{2}\|Ax - b\|^2$ for some matrix A and vector b . In this case, $L(\nabla f)$ is simply the spectral radius of $A^T A$.

For non-constant stepsizes, it is known that convergence is guaranteed if the stepsizes satisfy $0 < l < \tau^k < u < 2/L(\nabla f)$ for some upper bound u and lower bound l (see theorem 3.4 in [10] for this result and its generalization).

In practice, one seldom has accurate knowledge of $L(\nabla f)$, and the best stepsize choice depends on both the problem being solved and the error at each iteration. For this reason, it is better in practice to choose the sequence $\{\tau^k\}$ adaptively and enforce convergence using backtracking rules rather than the explicit stepsize restriction (6). These issues will be discussed in depth in Section 4.4.

3. APPLICATIONS OF FORWARD-BACKWARD SPLITTING

In this section we study a variety of problems, and discuss how they are formulated and solved using forward-backward splitting. FBS finds use in a large number of fields, including machine learning, signal and image processing, statistics, and communication systems. Here, we briefly discuss a small subset of potential applications. In Section 6 we present numerical experiments using these test problems.

3.1. Simple Applications.

3.1.1. *Convex Constraints and the Projected Gradient Method.* The projected gradient (PG) method is a special case of FBS involving a convex constraint. Suppose we are interested in solving

$$(7) \quad \text{minimize} \quad f(x) \quad \text{subject to} \quad x \in \mathcal{C}$$

for some convex set \mathcal{C} . We can rewrite this problem in the form (1) using the (non-smooth) characteristic function $\mathcal{X}_{\mathcal{C}}(x)$ of the set \mathcal{C} , which is zero for $x \in \mathcal{C}$ and infinity otherwise. As a consequence, the problem

$$\text{minimize} \quad f(x) + \mathcal{X}_{\mathcal{C}}(x)$$

is then equivalent to (1) with $g(x) = \mathcal{X}_{\mathcal{C}}(x)$. To apply FBS, we must evaluate the proximal operator of $\mathcal{X}_{\mathcal{C}}$:

$$(8) \quad \text{prox}_{\mathcal{X}_{\mathcal{C}}}(z, \tau) = \arg \min_{x \in \mathcal{C}} \frac{1}{2} \|x - z\|^2.$$

The solution to (8) is the element of \mathcal{C} closest to z ; this is simply the orthogonal projection of z onto the set \mathcal{C} . The resulting PG method was originally studied by Goldstein, Levitin, and Polyak [11, 12, 13].

3.1.2. *Lasso Regression.* One of the earliest sparse regression tools from high-dimensional statistics is the Lasso regression, which is easily written in the form (7).

An important application of PG methods as in (7) is the Lasso regression problem [14], which is an important sparse regression tool. The Lasso regression is defined as follows:

$$(9) \quad \text{minimize} \quad \frac{1}{2} \|Ax - b\|^2 \quad \text{subject to} \quad \|x\|_1 \leq \lambda.$$

Here, the convex set $\mathcal{C} = \{x : \|x\|_1 \leq \lambda\}$ is simply an ℓ_1 -norm ball. The corresponding proximal step, i.e., projection onto the ℓ_1 -norm ball, can be carried out efficiently using linear-time algorithms, such as the method proposed in [15]. We note that FBS has been used previously for the Lasso regression problem (9) in the popular SPGL1 solver [16, 17].

3.1.3. *ℓ_1 -Norm Penalized (Logistic) Regression.* One of the most common applications for FBS is the following ℓ_1 -norm penalized least squares problem:

$$(10) \quad \text{minimize} \quad \mu \|x\|_1 + \frac{1}{2} \|Ax - b\|^2.$$

In statistics, problem (10) is used to find sparse solutions to under-determined least squares problems. Problem (10) is called basis pursuit denoising (BDPN) in the context of compressive sensing and sparse signal recovery [8, 3].

Suppose the vector $b \in \{0, 1\}^M$ contains binary entries representing the outcomes of random Bernoulli trials. The success probability of the i th entry is $P(b_i = 1 | x) = e^{A_i x} / (1 + e^{A_i x})$ where A_i denotes the i th row of A . One is then interested in solving the so-called sparse *logistic* regression problem

$$(11) \quad \text{minimize} \quad \mu \|x\|_1 + \text{logit}(Ax, b)$$

with the logit penalty function defined as

$$\text{logit}(z, b) = \sum_{i=1}^M \log(e^{z_i} + 1) - b_i z_i.$$

Both problems, (10) and (11), can be solved using FBS with $g(x) = \|x\|_1$. The proximal operator of g is given by the well-known shrinkage operator $\text{shrink}(z, \mu\tau)$, whose i^{th} element is obtained as

$$(12) \quad \text{shrink}(z, \mu\tau)_i = \text{shrink}(z_i, \mu\tau) = \text{sign}(z_i) \max\{|z_i| - \mu\tau, 0\}.$$

3.1.4. Multiple Measurement Vector (MMV). In Section (3.1.3) we sought sparse vectors satisfying $Ax \approx b$ for some matrix A and measurement vector b . Suppose now that we have a matrix B containing many measurement vectors, and we want to find a sparse matrix X satisfying $AX \approx B$. If we further suppose that all columns of X must have the same sparsity pattern, then we arrive at the Multiple Measurement Vector (MMV) problem [18]. To formulate MMV using convex optimization, we need the group sparsity prior

$$\text{MMV}(X) = \sum_i \sqrt{\sum_j X_{ij}^2} = \sum_i \|X_i\|$$

where X_i denotes the i th row of X .

The convex formulation of MMV is

$$\underset{X}{\text{minimize}} \quad \mu \text{MMV}(X) + \frac{1}{2} \|AX - B\|^2$$

where the scalar μ controls the strength of the penalty. This is easily solved using FBS. The forward step is simply

$$\hat{X}^{k+1} = X^k - \tau^k A^T (AX^k - B).$$

The backward step requires the proximal operator of $\mu \text{MMV}(\cdot)$, which can be evaluated row-by-row. The i th row is simply

$$X_i^{k+1} = \text{prox}_{\mu \text{MMV}}(\hat{X}^{k+1}, \tau^k)_i = \hat{X}_i^{k+1} \frac{\max\{\|\hat{X}_i^{k+1}\| - \tau^k \mu, 0\}}{\|\hat{X}_i^{k+1}\|}.$$

3.1.5. Democratic Representations. The dynamic range of signals plays an important role in approximate nearest neighbor search and vector quantization [7], as well as communications and robotics/control [19]. Given a signal $b \in \mathbb{R}^M$, a low-dynamic range representation can be found by choosing a suitable matrix $A \in \mathbb{R}^{M \times N}$ with $M < N$, and by solving

$$(13) \quad \underset{x}{\text{minimize}} \quad \mu \|x\|_\infty + \frac{1}{2} \|Ax - b\|^2.$$

The problem (13) often yields a so-called democratic representation x^* that has small dynamic range and for which a large number of entries have equal (and small) magnitude [19]. The problem (13) can, once again, be solved using FBS. The only missing ingredient is the proximal operator for the ℓ_∞ -norm, which can be computed in linear time using the method in [15].

3.1.6. Low-Rank (1-bit) Matrix Completion. The goal of matrix completion is to recover a low-rank matrix \hat{X} from a small number of linear measurements. Such problems often involve the *nuclear norm* of a matrix, which is simply the sum of the magnitudes (i.e., the ℓ_1 norm) of the eigenvalues.

One prominent formulation of matrix completion takes the form

$$(14) \quad \underset{X}{\text{minimize}} \quad \mu \|X\|_* + L(X, Y),$$

where $\|X\|_*$ is the low-rank inducing nuclear norm of the matrix X and $L(\cdot, Y)$ is a loss function that depends on the model of the observed data Y [20].

Suppose that Y contains noisy and punctured (or missing) data of the entries of \hat{X} , and let Ω denote the set of indices that have been observed. In this case, one can use

$$L(X, Y) = \sum_{\alpha \in \Omega} (X_\alpha - Y_\alpha)^2$$

as an appropriate loss function. Another loss function arises in 1-bit matrix completion, which can be viewed as an instance of low-rank logistic regression [21]. Each observation Y_α is assumed to be a Bernoulli random variable with success probability $P(Y_\alpha = 1 | \hat{X}_\alpha) = e^{\hat{X}_\alpha} / (1 + e^{\hat{X}_\alpha})$. For this case, the appropriate loss function is again the logit function:

$$L(X, Y) = \text{logit}(X, Y).$$

For both quadratic and logit link functions, (14) can be solved using FBS. Specifically, one needs to compute the proximal operator of the nuclear norm, defined as

$$(15) \quad \text{prox}_*(Z, \tau) = \arg \min_X \tau \mu \|X\|_* + \frac{1}{2} \|X - Z\|^2,$$

whose solution is given by the matrix $U \text{shrink}(S, \mu\tau) V^T$, where $Z = USV^T$ is the singular value decomposition (SVD) of Z [21].

3.2. More Complex Applications. Sometimes an objective function does not immediately decompose into a smooth part and a “simple” part for which we can evaluate the proximal operator in closed form. In this case, it is often possible to re-formulate the problem in such a way that FBS is easily applied. In this section, we look at two such problems (total-variation denoising and support vector machines) that can be reformulated using *duality* and then solved easily and efficiently using FBS. We also look at an example where the *convex relaxation* of a problem is easily solved using FBS.

3.2.1. Total-Variation Denoising. Given a 2-dimensional image u , the intensity of the pixel in the i th row and j th column is denoted u_{ij} . The discrete gradient of u is denoted ∇u . At each point in the image, the gradient is given by $(\nabla u)_{ij} = (u_{i+1,j} - u_{i,j}, u_{i,j+1} - u_{i,j})^T$ and contains the discrete derivatives in the row and column directions.

Given a noise-contaminated image f , one can construct a denoised image by solving

$$(16) \quad \underset{u}{\text{minimize}} \quad \mu |\nabla u| + \frac{1}{2} \|u - f\|^2$$

where μ is a parameter to control the level of smoothing, and

$$(17) \quad |\nabla u| = \sum_{i,j} \|(\nabla u)_{ij}\| = \sum_{i,j} \sqrt{(u_{i+1,j} - u_{i,j})^2 + (u_{i,j+1} - u_{i,j})^2}$$

denotes the total-variation of u [22]. Rather than imposing sparsity on u itself, the total variation regularizer imposes sparsity on the *gradient* of u . As a result, problem (16) finds a piecewise-constant approximation to f .

To perform TV minimization using FBS, we must re-formulate $|\nabla u|$ into a simpler differentiable function. We begin by letting $x_{ij} = (x_{ij}^r, x_{ij}^c)^T$ be a vector in \mathbb{R}^2 . We can now write

$$(18) \quad \max_{\|x_{ij}\| \leq 1} \langle x_{ij}, (\nabla u)_{ij} \rangle = \|(\nabla u)_{ij}\|.$$

Equation (18) follows from the Cauchy-Swartz inequality, which states that $\langle x_{ij}, (\nabla u)_{ij} \rangle \leq \|x_{ij}\| \|(\nabla u)_{ij}\|$. Equality is attained by choosing x_{ij} parallel to $(\nabla u)_{ij}$. If we further

choose x_{ij} to have unit norm, then we arrive at (18). If we apply (18) to (17) we get $|\nabla u| = \max_{\|x\|_\infty \leq 1} \langle x, \nabla u \rangle$. We now have

$$(19) \quad \min_u \mu |\nabla u| + \frac{1}{2} \|u - f\|^2 = \min_u \max_{\|x\|_\infty \leq 1} \mu \langle x, \nabla u \rangle + \frac{1}{2} \|u - f\|^2 \\ = \max_{\|x\|_\infty \leq 1} \min_u \mu \langle x, \nabla u \rangle + \frac{1}{2} \|u - f\|^2.$$

The inner minimization in (19) is now differentiable. For a given x , the minimal value of u satisfies $u = f + \mu \nabla \cdot x$, where $\nabla \cdot x$ is the discrete divergence (which is the negative adjoint of the gradient operator). At pixel ij , this operator takes on the scalar value $(\nabla \cdot x)_{ij} = x_{i,j}^r - x_{i-1,j}^r + x_{i,j}^c - x_{i,j-1}^c$. If we plug the optimal value of u into (19) and simplify, we see that the optimal value of x is given by:

$$(20) \quad x^* = \arg \max_{\|x\|_\infty \leq 1} -\frac{1}{2} \|\nabla \cdot x + \frac{1}{\mu} f\|^2 = \arg \min_{\|x\|_\infty \leq 1} \frac{1}{2} \|\nabla \cdot x + \frac{1}{\mu} f\|^2.$$

This is simply a quadratic minimization with an infinity-norm constraint. Problem (20) is known as the *dual* form of (16). This problem can be solved using FBS as in Section (3.1.1). The resulting algorithm alternately performs gradient descent steps on (20), and then re-projects the result back into the infinity-norm ball using the formula $x_{ij} \leftarrow x_{ij} / \max\{\|x_{ij}\|, 1\}$. Once problem (20) has been solved, the optimal (denoised) image $u^* = f + \mu \nabla \cdot x^*$ is computed.

This approach to total-variation minimization was first taken in [23] using constant step-size parameters. A similar approach was taken using accelerated variants of FBS in [24].

3.2.2. Support Vector Machines. Consider a set of data points $\{d_i\}$ each of which lies in \mathbb{R}^n and has a binary label $l_i \in \{-1, 1\}$. The support vector machine (SVM) aims to find the hyper-plane in \mathbb{R}^n that separates the points with label +1 from the points with label -1. This is done by solving

$$(21) \quad \underset{w}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 + C \cdot h_l(Dw)$$

where $w \in \mathbb{R}^n$, C is a constant regularization parameter, D is a matrix with i th row equal to d_i , and h is the *hinge loss* function

$$h_l(z) = \sum_i \max\{1 - l_i z_i, 0\}.$$

The hinge loss function penalizes points that either lie on the “wrong” side of the hyper-plane $w^T x = 0$ or that lie too close to this hyperplane.

To apply FBS, we “dualize” this problem much like we did for total variation (Section 3.2.1). We first note that for any $z \in \mathbb{R}$, we have $C \max\{z, 0\} = \max_{0 \leq x \leq C} x z$. Using this observation, we re-write (21) as

$$(22) \quad \min_w \max_{0 \leq x \leq C} \frac{1}{2} \|w\|^2 + x^T (1_n - LDw) = \max_{0 \leq x \leq C} \min_w \frac{1}{2} \|w\|^2 + x^T (1_n - LDw)$$

where the constraints on the vector x are interpreted element-wise, L is a diagonal matrix with $L_{i,i} = l_i$, and 1_n is a vector of 1’s. This objective is quadratic in w , and the optimal value of w is given by $D^T L x$. If we plug this value in for w in (22) and multiply by -1 to convert the maximization for x into a minimization, we arrive at the dual problem

$$(23) \quad \underset{0 \leq x \leq C}{\text{minimize}} \quad \frac{1}{2} \|D^T L x\|^2 - \sum_i x_i.$$

This is a simple quadratic minimization with constraints of the form $0 \leq x_i \leq C$. The proximal operator corresponding to these constraints is the projection $\text{prox}(z, t) = \min\{\max\{z, 0\}, C\}$. Once the dual problem (23) is solved using FBS, the primal solution is recovered from the formula $w = D^T Lx$. SVM's and the dual minimization are reviewed in [25].

3.2.3. Phase Retrieval and Rank Minimization. A variety of non-convex problems can be relaxed into convex problems involving symmetric positive semi-definite (SPDP) matrices. This idea was popularized by [26], in which it was shown that the NP-complete MaxCut problem has a convex relaxation with the form of a semi-definite program. Here, we consider the more recent PhaseLift algorithm, which can be used for phase retrieval applications [27].

Suppose we wish to recover a complex-valued signal $x \in \mathbb{C}^N$ from measurements of the form $|\langle a_i, x \rangle|^2 = b_i$ for some set of vectors $\{a_i\}_{i=1}^M$. In words, we take inner products of the vector x with the measurement vectors $\{a_i\}$ and discard the phase information.

Recovery of x from the phase-less measurements b requires the solution of a system of non-linear equations. However, one can relax the equations into a convex problem by defining $A_i = a_i a_i^T$, and observing that $|\langle a_i, x \rangle|^2 = \langle A_i, x x^T \rangle = b_i$. Letting $\mathcal{A}(x x^T)_i = \langle A_i, x x^T \rangle$, the set of phase-less measurements can be written compactly as $\mathcal{A}(x x^T) = b$. Finally, note that any SPSPD rank-1 matrix X can be factorized as $X = x x^T$. This observation allows us to pose the recovery problem in the following form:

$$(24) \quad \text{minimize } \text{rank}(X) \quad \text{subject to } \mathcal{A}(X) = b, X \succeq 0.$$

Even though $\text{rank}(\cdot)$ is a non-convex function, one can form a convex relaxation of (24) by replacing $\text{rank}(\cdot)$ with the sparsity-inducing nuclear norm. The resulting problem, known as PhaseLift, is given by [27]:

$$(25) \quad \text{minimize } \|X\|_* \quad \text{subject to } \mathcal{A}(X) = b, X \succeq 0.$$

In the case where the measurement vector b is contaminated by additive noise, we choose an ℓ_2 -norm penalty, which stems from a Gaussian noise model. Specifically, we can write the resulting phase-retrieval problem as

$$\text{minimize } \mu \|X\|_* + \|\mathcal{A}(X) - b\|^2 \quad \text{subject to } X \succeq 0,$$

which can be solved using FBS. The relevant proximal operator corresponds to

$$\text{minimize } \tau \mu \|X\|_* + \|X - Z\|^2,$$

whose solution is given by the matrix U shrink($\Lambda, \mu\tau$) U^T , where $Z = U \Lambda U^T$ is the eigenvalue decomposition of Z .

3.3. Non-convex Problems. All applications discussed so far involved convex optimization. In practice, FBS also works quite well when applied to *non-convex* problems. However, in this case there are no theoretical guarantees that the method will converge to a global minimizer, or that the algorithm will even converge at all. That being said, it is often the case for non-convex problems that no known method can guarantee optimality in polynomial time, and so this lack of theoretical guarantees makes FBS no weaker than other options.

When solving non-convex problems the user should be cautious of several things. First, unlike in the convex case, the final result will be sensitive to the initial iterate. For this reason, it is important to either initialize the methods with a good approximate solution, or else generate numerous solutions from different random initializations and choose the solution with the lowest objective value. Second, because the method is no longer guaranteed to converge, the user must put an intelligent limit on the number of iterations. In some

situations, it may even be best to run the algorithm for a pre-determined number of iterations rather than using one of the precision-based termination conditions to be discussed in Section 4.6.

3.3.1. Non-Negative Matrix Factorization. Many problems in information science require a matrix to be factored into low-rank components. Given a matrix $Q \in \mathbb{R}^{M \times N}$, we seek $W \in \mathbb{R}^{M \times K}$ and $C \in \mathbb{R}^{N \times K}$ with

$$Q \approx WC^T.$$

In *latent factor analysis*, the matrices W and C represent the underlying factors that “explain” the data. Consider the performance of students on exam questions, each of which requires a blend of knowledge from K different topics. Suppose the vector $w_p \in \mathbb{R}^K$ measures how much of each topic is required to solve problem p , and $c_s \in \mathbb{R}^K$ describes a student’s mastery of each of the K topics. If we choose these vectors appropriately, we expect student s to score $Q_{ps} = \langle w_p, c_s \rangle$ on problem p . Intuitively, it seems that a student’s skill level should always be greater than or equal to zero, and each problem should require a non-negative amount of knowledge from each topic. For this reason, it makes sense to recover the difficulty and skill vectors by solving

$$(26) \quad \begin{aligned} & \underset{W, C}{\text{minimize}} \quad \|Q - WC^T\|^2 \\ & \text{subject to} \quad W \geq 0, C \geq 0 \end{aligned}$$

where the constraints on W and C are interpreted element-wise.

This problem is known as *non-negative matrix factorization* (NMF) [28]. While various methods have been proposed to solve this problem [29, 30], forward backward splitting remains one of the simplest. The gradient of (26) is easily obtained using the chain rule, resulting in the forward step

$$\begin{pmatrix} \widehat{W}^{k+1} \\ \widehat{C}^{k+1} \end{pmatrix} = \begin{pmatrix} W^k \\ C^k \end{pmatrix} - \tau^k \begin{pmatrix} (W^k (C^k)^T - Q) C^k \\ (W^k (C^k)^T - Q)^T W^k \end{pmatrix}.$$

The backward step is a trivial projection onto the set of non-negative matrices – i.e., negative entries in the matrices are replaced by zeros.

3.3.2. Max-Norm Optimization and the Max Cut Problem. Any $N \times N$ symmetric positive semidefinite matrix A with rank at most K can be written in the form $A = XX^T$ for some matrix $X \in \mathbb{R}^{N \times K}$. Given such a representation, the max-norm of A is given by

$$\|A\|_{\max} = \max_i \|X_i\|_2^2 := \|X\|_{2, \infty}^2$$

where X_i is the i th row of X . It can be shown that the max-norm of a matrix is a good approximation to the nuclear norm (up to a constant factor), and can be used to promote low-rank solutions when the nuclear norm is not available because of the expense of computing singular value decompositions [31].

An important application of the max-norm is for solving the max-cut problem on a graph [31]. Consider a graph with N vertices and an edge of weight w_{ij} between vertex i and j . An assignment of a binary label $x_i \in \{-1, 1\}$ to each vertex is called a “cut.” The “value” of a cut is the sum of all edge weights connecting vertices with different labels.

Finding the graph cut with maximum value is NP-complete [26], however a good approximation can be found by solving

$$(27) \quad \begin{aligned} & \underset{X}{\text{minimize}} \quad \langle W, XX^T \rangle \\ & \text{subject to} \quad \|X\|_{2,\infty}^2 \leq 1 \end{aligned}$$

where W is the weighted adjacency matrix of the graph, and $X \in \mathbb{R}^{N \times K}$ is a matrix [31]. Once a solution to (27) is found, a labeling is generated by computing $x_i = \text{sign}(X_i s)$ for some random vector s .

The objective in (27) differentiable, and the forward step of FBS is given by

$$\hat{X}^{k+1} = X^k - \tau^k X(W^T + W).$$

The backward step is simply a projection onto the constraint set. In this case, we have

$$X_i^{k+1} = \hat{X}_i^{k+1} \frac{\min\{\|\hat{X}_i^{k+1}\|, 1\}}{\|\hat{X}_i^{k+1}\|}.$$

4. BELLS AND WHISTLES

FBS in its raw form suffers from numerous problems including: (i) The convergence speed of FBS depends strongly on the choice of the stepsize parameters $\{\tau_k\}$. The best stepsize choices may not be intuitively obvious. (ii) Convergence is only guaranteed when the stepsizes satisfy the stability condition (6) which depends on the Lipschitz constant for ∇f . In real applications one often has no knowledge of this Lipschitz constant and no practical way to estimate it. (iii) To automatically stop the FBS iteration, a good measure of convergence is needed.

In this section, we discuss practical methods for overcoming these problems. We begin with adaptive schemes for automatic stepsize selection. We then discuss backtracking methods that can guarantee stability even when the user has no explicit knowledge of $L(\nabla f)$. Finally, we discuss stopping conditions and measures of convergence.

4.1. Adaptive Stepsize Selection. The efficiency of FBS (and gradient methods in general) is very sensitive to the choice of the stepsize τ_k . For this reason, much work has been devoted to studying *adaptive* stepsize strategies. Adaptive methods automatically tune stepsize parameters in real time (as the algorithm runs) to achieve fast convergence. In this section, we discuss spectral (also called Barzilai-Borwein) stepsize methods, and how they can be adapted to FBS.

Before considering the full-scale FBS method for (1), we begin by considering the case $g = 0$. In this case $h = f$ and FBS reduces to simple gradient descent of the form

$$(28) \quad x^{k+1} = x^k - \tau^k \nabla f(x).$$

Spectral schemes for gradient descent were proposed by Barzilai and Borwein [32], who model the function f as the simple quadratic function

$$(29) \quad f(x) \approx \hat{f}(x) = \frac{a}{2} \|x\|^2 + \langle x, b \rangle.$$

It can be shown that the optimal stepsize choice for (29) is $\tau = 1/\alpha$. With this choice, the gradient descent method achieves a perfect minimizer of the simple quadratic (29) in one iteration. This motivates the following stepsize scheme for gradient descent: Before applying the descent step (28), approximate f with a quadratic of the form (29). Then, take a step of length $\tau^k = 1/\alpha$.

Spectral stepsize rules have been generalized to handle FBS with specific choices of g . When g is the characteristic function of a convex set, the marriage of FBS with an adaptive stepsize produces the *spectral projected gradient method* (SPG) [33]. SPG was further specialized to handle the case of ℓ_1 constraints by van den Berg and Friedlander [16, 17]. The authors of [34] present a spectral stepsize rule that handles the case $g = |\cdot|$. The resulting method solves ℓ_1 penalized least squares problems using a combination of spectral gradient descent and subspace optimization. This idea was revisited in [35], in which the authors present a modification of FBS for ℓ_1 penalized least squares that implicitly performs conjugate gradient optimization inside of low-rank subspaces.

Spectral stepsizes for general FBS with arbitrary f and g were first studied in [36] for use with the solver SpaRSA. The method presented here is a hybrid of the spectral scheme [36] with the “adaptive” stepsize rule presented in [37].

It was observed in the introduction that the stepsize restriction for FBS depends only on f and not on g . Spectral methods for FBS exploit this property. The idea is to build a quadratic approximation for f of the form (29) at each iteration, and then choose the optimal gradient descent stepsize $\tau^k = 1/a$. Let

$$(30) \quad \Delta x^k = x^k - x^{k-1},$$

$$(31) \quad \Delta F^k = \nabla f(x^k) - \nabla f(x^{k-1}).$$

If we assume a quadratic model for f of the form (29), then we have

$$\Delta F^k = \nabla f(x^k) - \nabla f(x^{k-1}) = a(x^k - x^{k-1}).$$

The function (29) is fit to f using a least squares method which chooses a to minimize either $\|\Delta F^k - a\Delta x^k\|^2$ or $\|a^{-1}\Delta F^k - \Delta x^k\|^2$. Just like in the case $g = 0$, we then select a stepsize of length $\tau^k = 1/a$. The resulting stepsize choice is given by

$$(32) \quad \tau_s^k = \frac{\langle \Delta x^k, \Delta x^k \rangle}{\langle \Delta x^k, \Delta F^k \rangle}, \text{ and } \tau_m^k = \frac{\langle \Delta x^k, \Delta F^k \rangle}{\langle \Delta F^k, \Delta F^k \rangle}$$

respectively for each least squares problem. The value τ_s^k is known as the “steepest descent” stepsize, and τ_m^k is called the “minimum residual” stepsize [37].

A number of variations of spectral descent are reviewed by Fletcher [38], several of which perform better in practice than the “classic” stepsize rules (32). We recommend the “adaptive” BB method [37], which uses the rule

$$(33) \quad \tau^k = \begin{cases} \tau_m^k & \text{if } \tau_m^k / \tau_s^k > \frac{1}{2} \\ \tau_s^k - \frac{1}{2}\tau_m^k & \text{otherwise.} \end{cases}$$

Note that (particularly for non-convex problems), the stepsize τ_m^k or τ_s^k may be negative. If this happens, the stepsize should be discarded and replaced with its value from the previous iteration. When complex-valued problems are considered, it is important to use only the real part of the inner product in (33).

4.2. Preconditioning. FBS can be “preconditioned” by multiplying the gradient directions of the forward step with a symmetric positive definite matrix Γ . The backward step must be modified as well to ensure convergence. The resulting approach, sometimes called the *proximal newton method* [39, 40], is described in Algorithm 2.

Algorithm 2 Preconditioned Forward-Backward Splitting

```

while not converged do
(34)    $\hat{x}^{k+1} = x^k - \tau^k \Gamma \nabla f(x^k)$ 
(35)    $x^{k+1} = \text{prox}_g(\hat{x}^{k+1}, \tau^k, \Gamma) = \arg \min_x \tau^k g(x) + \frac{1}{2} \|x - \hat{x}^{k+1}\|_{\Gamma^{-1}}^2$ 
end while

```

The preconditioned backward step (35) now requires the *generalized* proximal operator, which involves the following norm weighted by the matrix Γ^{-1}

$$\|z\|_{\Gamma^{-1}}^2 = \langle \Gamma^{-1} z, z \rangle.$$

For general Γ the minimization (35) may not have a closed form solution, even when g is “simple.” However, when the function g acts separately on each element of x , adding a *diagonal* preconditioner is a trivial modification. For example, when $g = \mu|\cdot|$ (i.e., for the sparse least squares problems in section 3.1.3) and Γ is diagonal, the i th entry of $\text{prox}_g(z, \tau, \Gamma)$ is simply given by $\text{shrink}(z_i, \mu\tau\Gamma_{ii})$. Preconditioning with a diagonal Γ is similarly easy in the case of the Support Vector Machine (Section 3.2.2).

The preconditioner has another interpretation. One can derive the preconditioned iteration (Algorithm 2) by making the change of variables $x \rightarrow Py$ for some invertible matrix P and solving the problem

$$(36) \quad \underset{y}{\text{minimize}} \quad f(Py) + g(Py).$$

Each iterate y^k can then be converted to an approximate solution to (1) by computing $x^k = Py^k$. It can be shown that the iterates $\{x^k\}$ obtained through this procedure are equivalent to those obtained by Algorithm 2 with preconditioner $\Gamma = P^T P$.

The latter interpretation of the preconditioner is important because it suggests how to choose Γ . When f has the form

$$f = \hat{f}(Ax)$$

for some matrix A and function \hat{f} , the preconditioned problem (36) involves the objective $\hat{f}(APy)$. For example, when $f(x) = \|Ax - b\|^2$, the problem (36) contains the term $\|APy - b\|^2$. If A is poorly conditioned, then one should choose P to be a *diagonal* matrix such that AP is more well conditioned. One effective method chooses the entries in P so that all columns of AP have unit norm. The corresponding choice of $\Gamma = P^T P$ is then

$$\Gamma_{ii} = \frac{1}{\|a_i\|^2}$$

where a_i denotes the i th column of A .

Note that, because of the equivalence between preconditioned FBS and the original FBS on problem (36), all the tricks described in this section (including adaptivity and backtracking) can be applied to Algorithm 2 as long as Γ remains unchanged between iterations. In the event that Γ changes on every iteration, the converge theory becomes somewhat more complex (see [9]).

4.3. Acceleration. Adaptive stepsize selection helps to speed the convergence of the potentially slow FBS method. Another approach to dealing with slow convergence is to use predictor-corrector schemes that “accelerate” the convergence of FBS.

While several predictor-corrector variants of FBS have been proposed (see for example [41]), the algorithm FISTA has become quite popular because of its lack of tuning parameters and good worst-case performance. FISTA relies on a one-size-fits-all sequence of acceleration parameters that works for any objective. FISTA is listed in Algorithm 3. Step 2 of FISTA simply performs an FBS step. Step 4 performs *prediction*, in which the current iterate is advanced further in the direction it moved during the previous iteration. The aggressiveness of this prediction step is controlled by the scalar parameter α^k . This acceleration parameter is updated in step 3, and increases on each iteration causing the algorithm to become progressively more aggressive. The FISTA method alternates between “predicting” an aggressive estimate of the solution, and “correcting” this estimate using FBS to attain greater accuracy.

One notable advantage of this approach is the worst-case convergence rate. It has been shown that FISTA decreases the optimality gap (the difference between the objective at iteration k and the optimal objective value) with rate $O(\frac{1}{k^2})$ [42]. In contrast, the worst-case performance of conventional FBS is known to be $O(\frac{1}{k})$.

Algorithm 3 FISTA

Require: $y^1 = x^0 \in \mathbb{R}^N$, $\alpha^1 = 1$, $\tau < 1/L(\nabla G)$

```

1: for  $k = 1, 2, 3, \dots$  do
2:    $x^k = \text{prox}_g(y^k - \tau \nabla f(y^k), \tau)$ 
3:    $\alpha^{k+1} = (1 + \sqrt{1 + 4(\alpha^k)^2})/2$ 
4:    $y^{k+1} = x^k + \frac{\alpha^k - 1}{\alpha^{k+1}}(x^k - x^{k-1})$ 
5: end for
```

Rather than requiring the user to choose a stepsize, FISTA is generally coupled with a backtracking line search such as those described in the next section.

In theory, FISTA achieves superior worst-case performance by increasing the prediction parameter α^k on every iteration. In practice, however, FISTA does not always perform well when the prediction parameter α^k becomes large. This is particularly true when a large/aggressive stepsize is used, in which case the objective values may oscillate rapidly before convergence is reached. For this reason FISTA (and other Nesterov-type methods) generally perform better when a “restart” method is used [43]. Such methods allow α^k to increase while the performance of Algorithm 3 is good, but reset the value to $\alpha^k = 1$ if oscillations develop at iteration k .

One restart strategy simply sets $\alpha^k = 1$ on iterations where the objective value increases. However, this method requires the objective value to be calculated at each iteration, which is expensive and in general unnecessary. The authors of [43] suggest a restart method that does not require the objective. The method restarts when the difference between iterates $x^k - x^{k-1}$ points in an ascent direction for the objective. It can be shown that this occurs whenever

$$(37) \quad \langle y^k - x^k, x^k - x^{k-1} \rangle \geq 0$$

in which case the method should be restarted [43]. We will see in Section 4.6 that the vector $y^k - x^k$ is parallel to the gradient of the objective function. For this reason, condition (37) restarts the method whenever the change between iterates forms an acute angle with the gradient (which is the direction of steepest ascent).

Algorithm 4 Non-Monotone Line Search

```

while  $x^k$  and  $x^{k+1}$  violate condition (38) do
     $\tau^k \leftarrow \tau^k / 2$ 
     $x^{k+1} \leftarrow \text{prox}_g(x^k - \tau \nabla f(x^k), \tau)$ 
end while

```

4.4. Backtracking Line Search. Convergence of FBS can be guaranteed by enforcing the stability condition (6). In practice, though, the user generally has no knowledge of the global properties of ∇f , and so the actual value of this stepsize restriction is unknown. In Section 4.1 we discuss adaptive methods that automatically choose stepsizes for us. This does not free us from the need for stability conditions: spectral methods are, in general, not guaranteed to converge, even for convex problems [33].

Even without knowing the stepsize restriction (6), convergence can be guaranteed by incorporating a *backtracking line search*. Such methods proceed by checking a line search condition after each iteration of FBS. The search condition usually enforces that the objective has decreased sufficiently. If this condition fails to hold, then *backtracking* is performed – the stepsize is decreased and the FBS iteration repeated until the backtracking condition (i.e., sufficient decrease of the objective) is satisfied.

Line search methods were originally proposed for smooth problems [44], and later for projected gradient schemes [45, 33]. A backtracking scheme for general FBS was proposed by Beck and Teboulle for use with FISTA [42]. Many authors consider these line searches to be overly conservative, particularly for poorly conditioned problems [33], and prefer non-monotone line search conditions [46, 47]. Rather than insisting upon an objective decrease on every iteration, non-monotone line search conditions allow the objective to increase within limitations.

Non-monotone line search methods are advantageous for two reasons. First, for poorly conditioned problems (i.e., objective functions with long, narrow valleys) it may be the case that iterate x^{k+1} lies much closer to the minimizer than x^k , despite x^{k+1} having the larger objective value. A non-monotone line search prevents such iterates from being rejected.

Second, the objective function must be evaluated every time the line-search condition is tested. For complex problems where evaluating the objective is costly and multiple backtracking steps are necessary for each iteration, the line search procedure may dominate the runtime of the method. Non-monotone line search conditions are less likely to be violated and thus backtracking terminates faster, which alleviates this computational burden.

The method proposed here is inspired by the monotone search proposed in [42] for general FBS, and generalizes the non-monotone strategy of [46] for SPG.

Let $M > 0$ be an integer line search parameter, and define

$$\hat{f}^k = \max\{f^{k-1}, f^{k-2}, \dots, f^{k-\min\{M, k\}}\}.$$

After each step of FBS, the following line search condition is checked.

$$(38) \quad f(x^{k+1}) < \hat{f}^k + \Re\langle x^{k+1} - x^k, \nabla f(x^k) \rangle + \frac{1}{2\tau^k} \|x^{k+1} - x^k\|^2.$$

If the backtracking condition (38) fails, the stepsize is decreased until (38) is satisfied. This process is formalized in Algorithm 4. Note that Algorithm 4 always terminates because condition (38) is guaranteed to hold whenever τ^k is less than the reciprocal of the Lipschitz constant of ∇f .

A convergence proof for FBS with the line search described in Algorithm 4 is given in the appendix.

4.5. Continuation. Some types of regression problems are solved very quickly when solutions are highly sparse, but become numerically difficult when solutions lack sparsity. This is sometimes the case for problem (10). When μ is large, solutions to (10) are highly sparse. For small μ , solutions may lack sparsity, resulting in slow convergence if A is poorly conditioned.

Continuation methods for (10) exploit this observation by choosing a large initial value of μ and decreasing this parameter over time. This way, the solver can use the results of “easy” problems as a warm start for more difficult problems with less sparsity. Furthermore, using a warm start keeps the support of each iterate small, whereas the support might “blow up” for many iterations if difficult problems are attacked with a bad initializer.

Continuation techniques have been proposed by a number of authors (see for example [8, 48]) and are a keystone component in methods such as fixed-point continuation [49] that aim to approximate solutions to under-determined problems of the form

$$(39) \quad \text{minimize} \quad \mu \|x\|_1 \quad \text{subject to} \quad Ax - b = 0$$

by solving (10) with very small μ .

When choosing values for μ , it helps to observe that the solution to (10) is zero when $\mu \geq \|A^T b\|_\infty$. For this reason it is suggested to choose an initial parameter of $\mu = \eta \|A^T b\|_\infty$ for some $\eta < 1$. After we obtain a solution using the current value of μ , we replace $\mu \leftarrow \eta \mu$ and continue iterating until the desired value of μ is reached. We suggest to choose $\eta = \frac{1}{5}$, although the practical performance of continuation is not highly sensitive to this parameter.

Continuation is mostly effective for problem (10) when A is poorly conditioned, or when extremely large regularization parameters are needed. Note that continuation does not always enhance performance, and may sometimes even make performance dramatically worse.

4.6. Stopping Conditions. While the accuracy of FBS becomes arbitrarily good as the number of iterations approaches infinity, we must of course stop after a finite number of iterations. A good stopping criteria should be strict enough to guarantee an acceptable degree of accuracy without requiring an excessive number of iterations.

Our stopping conditions will be based on the *residual*, which is simply the derivative of the objective function (or a sub-gradient in the case that g is non-differentiable). Because f is assumed to be smooth, we can differentiate this term in the objective directly. While we may not be able to differentiate g , we see from equation (5) that a sub-gradient is given by $(\hat{x}^{k+1} - x^{k+1})/\tau^k \in \partial g(x^{k+1})$. We now have the following formula for the residual r^{k+1} at iterate x^{k+1} :

$$(40) \quad r^{k+1} = \nabla f(x^{k+1}) + \frac{\hat{x}^{k+1} - x^{k+1}}{\tau^k}.$$

A simple termination rule would stop the algorithm when $\|r^{k+1}\| < \text{tol}$ for some small tolerance $\text{tol} > 0$. However, this rule is problematic because it is not *scale invariant*. To understand what this means, consider the minimization of some objective function $h(\cdot)$. This function can be re-scaled by a factor of 1000 to obtain $\hat{h} = 1000h$. The new rescaled objective has the same minimizer as the original, however the sub-gradient $\partial \hat{h}(x^k)$ is 1000 times larger than $\partial h(x^k)$. A stopping parameter tol may be reasonable for minimizing h but overly strict for minimizing \hat{h} , even though the solutions to the problems are identical. Ideally, we would like scale invariant stopping rules that treat both of these problems equally.

One way to achieve scale invariance is by replacing the residual with the *relative residual*. To define the relative residual, we observe that (40) is small when

$$(41) \quad \nabla f(x^{k+1}) \approx -\frac{\hat{x}^{k+1} - x^{k+1}}{\tau^k}.$$

In plain words, the residual (40) measures the difference between the gradient of f and the negative sub-gradient of g . The relative residual r_r^{k+1} measures the *relative* difference between these two quantities, which is given by

$$(42) \quad r_r^{k+1} = \frac{\|\nabla f(x^{k+1}) + \frac{\hat{x}^{k+1} - x^{k+1}}{\tau^k}\|}{\max\{\|\nabla f(x^{k+1})\|, \|\frac{\hat{x}^{k+1} - x^{k+1}}{\tau^k}\|\} + \epsilon_r} = \frac{\|r^{k+1}\|}{\max\{\|\nabla f(x^{k+1})\|, \|\frac{\hat{x}^{k+1} - x^{k+1}}{\tau^k}\|\} + \epsilon_r}$$

where ϵ^r is some small positive constant to avoid dividing by zero.

Another more general scale invariant stopping condition uses the normalized residual, which is given by

$$(43) \quad r_n^{k+1} = \frac{\|r^{k+1}\|}{\|r^1\| + \epsilon_n}$$

where the small constant ϵ_n prevents division by zero. Rather than being an absolute measure of accuracy, the normalized residual measures how much the approximate solution has improved relative to x^1 .

Both scale invariant conditions have advantages and disadvantages. The relative residual works well for a wide range of problems and is insensitive to the initial choice of x^0 . However, the relative residual loses scale invariance when $\nabla f(x^*) = 0$ (in which case the denominator of (42) nearly vanishes). This happens, for example, when g is the characteristic function of a convex set and the constraint is inactive at the optimal point (i.e., the optimal point lies in the interior of the constraint set). In contrast, the normalized residual can be effective even if $\nabla f(x^*) = 0$. However, this measure of convergence is potentially sensitive to the choice of the initial iterate, and so it requires the algorithm to be initialized in some consistent way. The strictness of this condition can also depend on the problem being solved.

For general applications, we suggest a combined stopping condition that terminates the algorithm when either r_r^{k+1} or r_n^{k+1} gets small. In this case, r_r^{k+1} terminates the iteration if a high degree of accuracy is attained, and r_n^{k+1} terminates the residual appropriately in cases where $0 \in \partial g(x^*)$.

5. FASTA: A HANDY FORWARD-BACKWARD SOLVER

To create a common interface for testing different FBS variants, we have created the solver FASTA (Fast Adaptive Shrinkage/Thresholding), which implements forward-backward splitting for arbitrary problems. Many improvements to FBS are implemented in FASTA including adaptively, acceleration, backtracking, a variety of stopping conditions, and more. FASTA enables different FBS variants to be compared objectively while controlling for the effects of stepsize rules, programming language, and other implementation details.

Rather than addressing the problem (1) directly, FASTA solves general problems of the form

$$(44) \quad \text{minimize} \quad h(x) = \tilde{f}(Ax) + g(x),$$

where \tilde{f} and A are chosen so that $\tilde{f}(Ax) = f(x)$. This problem form allows FASTA to compute Ax^k once per iteration, and use the result twice to evaluate both the objective (which requires $\tilde{f}(Ax)$), and its gradient (given by $A^T \nabla \tilde{f}(Ax)$). The user can solve an arbitrary

problem by supplying A , the gradient of f and the proximal mapping (i.e., backward gradient descent operator) for g . However, custom wrappers are provided that solve all of the test problems described in this article, and codes are provided to reproduce the results in Section 6.

6. NUMERICAL EXPERIMENTS

We compare several variants of FBS and compare their performance using the test problems described in Section 3. The variants considered here are the original FBS with constant stepsize, and the accelerated variant FISTA described in Section 4.3. We also consider the adaptive method described in Section 4.1, which is implemented in the solver FASTA. This method uses the spectral stepsize rules proposed in [36] and [37].

The stepsize parameter for all methods was initialized by first estimating the Lipschitz constant $L(\nabla f)$. Two random vectors x^1, x^2 were generated, and the estimate $\tilde{L} = \|\nabla f(x^2) - \nabla f(x^1)\| / \|x^2 - x^1\| \leq L(\nabla f)$ was formed. The initial stepsize was then $\tau^0 = 10/\tilde{L}$. This stepsize is guaranteed to be at least an order of magnitude larger than the true stepsize restriction for FISTA ($1/L$). The backtracking scheme in Section 4.4 was then used to guarantee convergence. We found that this implementation enables faster convergence than implementations that explicitly require a Lipschitz constant for ∇f .

At each iteration of the algorithms considered, the relative residual (42) was computed and used as a measure of convergence. Time trials for all methods were terminated on the first iteration that satisfied $r_r^{k+1} < 10^{-4}$. If this stopping condition was not reached, iterations were terminated after 1000 iterations (except for the SVM problem, which was allowed up to 5000 iterations). Time trial results are averaged over 100 random trials.

6.1. Test Problem Details.

Lasso: We generated problems of the form (9) using random Gaussian matrices $A \in \mathbb{R}^{M \times N}$ with $N = 1000$. Experiments were done with both $M = 500$ (moderately under-sampled) and $M = 100$ (highly under-sampled). The true signal x^0 had 20 entries of unit magnitude; the remaining entries were zero. The indices of the support set were chosen at random. The measurement vector was generated using the formula $b = Ax^0$ and then contaminated by additive Gaussian noise to achieve a signal-to-noise ratio (SNR) of 13 dB. Recovery was performed with $\lambda = 15$.

ℓ_1 -Norm Penalized Least Squares: We generated problems of the form (10) using a similar procedure as for Lasso. This time, however, with the noise scaled to achieve SNR = 20 dB. Recovery was performed with $\mu = 0.1$.

Logistic Regression: We generated random matrices $A \in \mathbb{R}^{M \times N}$ of dimension $N = 1000$ and $M = 500$ with zero-mean Gaussian entries of variance of 4. Once again, the support of the true signal comprised 20 randomly chosen elements. Given $z = Ax$, b was constructed as a set of random Bernoulli variables with probability of success given by the logistic function $P(b_i = 1) = e^{z_i} / (1 + e^{z_i})$. The problem (11) was solved with $\mu = 20$.

Multiple Measurement Vector: We use the synthetic test problem suggested by [18]. A random Gaussian measurement matrix $A \in \mathbb{R}^{20 \times 30}$ was selected for each problem instance. A true solution matrix $X_0 \in \mathbb{R}^{30 \times 10}$ was chosen with 7 non-zero rows and random Gaussian entries in each non-zero row. The data matrix was $B = AX_0 + \eta$, where η was a random Gaussian noise matrix with $\sigma = 0.1$. The solution was recovered with $\mu = 1$.

Democratic Representations: We generated frames of dimension 500×1000 by randomly selecting subset of rows from a unitary discrete Fourier transform matrix. The signal b was

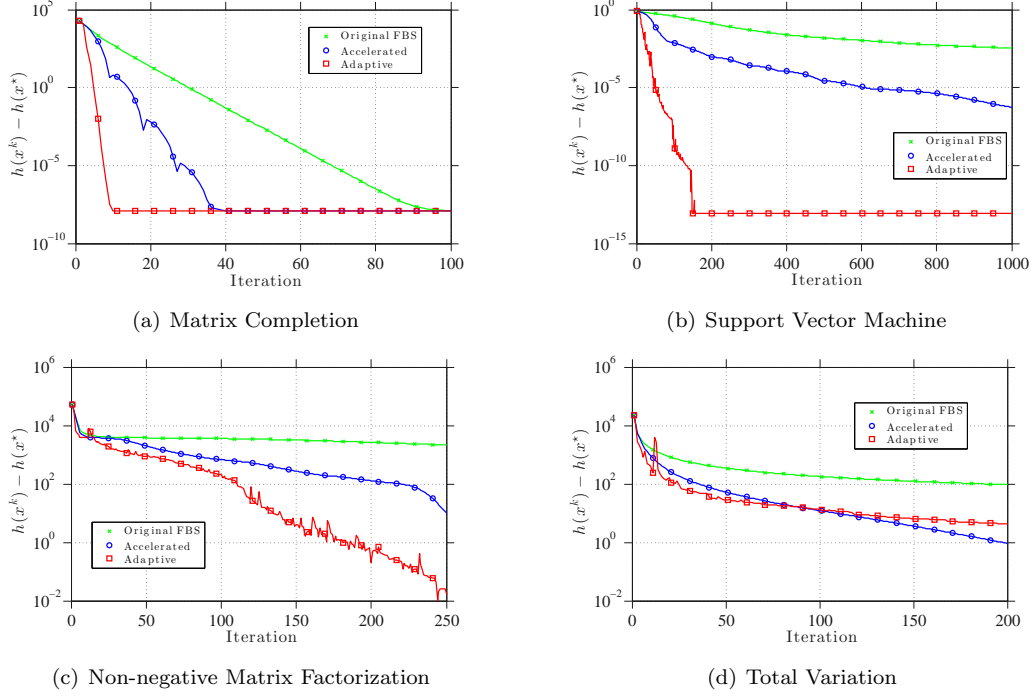


FIGURE 1. Sample converge curves for FBS, FBS+acceleration (FISTA), and FBS+adaptivity (SpARSA) for four diverse test problems (see Section 3 and 6 for details). The vertical axis shows the optimality gap.

a complex-valued random Gaussian vector of length 500. Equation (13) was solved with $\mu = 300$.

Matrix Completion: We generated a 200×1000 matrix X with i.i.d. random Gaussian entries of standard deviation 10. The SVD of X was then computed, and all but the largest 5 singular values were set to zero to obtain a rank-5 matrix. The sigmoidal logistic function was then applied element-wise to X to obtain a matrix of probabilities; a random Bernoulli matrix Y was drawn from the resulting distribution. Problem (14) was then solved using the logistic penalty function with $\mu = 25$.

Total Variation Denoising: The 256×256 Shepp-Logan phantom was constructed with pixel intensities that spanned the unit interval. Test images were then contaminated with Gaussian noise ($\sigma = 0.05$). Images were denoised by solving (16) with $\mu = 0.1$.

Support Vector Machine: Feature vectors were generated from two classes: one class containing random Gaussian variables with mean -1 , and one with mean $+1$. Each trial contained 1000 feature vectors of dimension 15. A support vector machine was trained to separate the two classes with regularization parameter $C = 10^{-2}$.

Phase Retrieval: We generated random vectors of length $N = 200$ with random Gaussian real and imaginary parts. The set $\{a_i\}$ was created by randomly drawing complex Gaussian vectors. The measurement vector b of length $M = 600$ was created by setting $b_i = |\langle a_i, x \rangle|^2$, and then, contaminating the resulting measurements with real-valued Gaussian noise to have $\text{SNR} = 13 \text{ dB}$. Problem (26) was then solved with parameter $\mu = 15$, which was chosen to be large enough that the solution matrix had rank 1.

TABLE 1. Complexity comparison of “vanilla” FBS, accelerated FBS (FISTA), and adaptive FBS (SpaRSA). We show the average number of iterations (and time per problem in seconds) to reach convergence. Adaptive FBS clearly outperforms other FBS variants for the considered test problems.

Problem	FBS	Accelerated	Adaptive
Lasso 100	356 (0.138)	55 (0.045)	22 (0.021)
Lasso 500	47 (0.065)	20 (0.039)	8 (0.018)
BPDN 100	253 (0.088)	48 (0.036)	20 (0.022)
BPDN 500	67 (0.077)	23 (0.039)	10 (0.019)
Logistic	40 (0.140)	24 (0.091)	14 (0.057)
MMV	657 (0.193)	81 (0.046)	58 (0.037)
Democratic	71 (0.074)	31 (0.044)	12 (0.028)
Mat Comp	69 (6.155)	26 (2.569)	8 (0.750)
TV Denoising	1000 (7.877)	177 (1.800)	102 (1.186)
SVM	3081 (0.643)	244 (0.090)	36 (0.024)
PhaseLift	1000 (36.720)	186 (6.967)	83 (3.614)
NMF	1000 (3.541)	246 (1.319)	173 (0.738)
Max-Norm	181 (16.578)	43 (5.939)	10 (0.909)

Non-Negative Matrix Factorization: A rank-10 factorization problem was built from two matrices $X \in \mathbb{R}^{800 \times 10}$ and $Y \in \mathbb{R}^{200 \times 10}$ consisting of uniform random numbers from the interval $[0, 1]$. The matrix $Q = XY^T$ was formed, and then contaminated with Gaussian noise of variance 10^{-2} . The matrix was then recovered by solving (26).

Max-Norm Optimization: We solve the max-cut graph segmentation problem described in [31]. A “two-moons” data set was built with 1000 data points $\{x_i\}$, and the weighted adjacency matrix $W \in \mathbb{R}^{1000 \times 1000}$ was built with $W_{ij} = \delta - e^{\|x_i - x_j\|^2 / \sigma^2}$, where $\delta = 0.01$ is a regularization parameter, and $\sigma = 0.1$ is a scaling parameter. Problem (27) was solved using the weighted adjacency matrix. The recovered matrix X was then used to compute the labeling given by $\text{sign}(Xr)$ where r is a random Gaussian vector.

6.2. Results and Discussion. For all problems considered, both the accelerated and adaptive FBS dramatically outperformed “vanilla” FBS without these modifications. However FBS does not reliably converge when adaptivity and acceleration are used simultaneously, and so the user must pick one.

To explore the efficiency of different approaches, we applied three variants of FBS to each test problem: Plain FBS, FBS with acceleration (FISTA), and FBS with adaptive stepsizes (SpaRSA). The accelerated FBS was accompanied by the restart rule (37), and all methods used backtracking line search. For each algorithm, the number of iterations (and time in seconds) needed to solve each test problem is reported in Table 1.

For most problems considered, the adaptive method outperformed the accelerated scheme by a factor of 3-to-5. For the SVM problem, the performance gap was somewhat larger. We can take a closer look at the behavior of each method with the convergence curves in Figure 1. Figure 1a shows the convergence for matrix completion, which looks typical for most problems including Lasso, penalized least-squares, logistic regression, and MMV. For such problems, we see smooth, exponential decay of the error until machine precision

is reached. Note this empirical behavior is much better than the $O(1/k)$ worse-case global convergence bounds [42].

We also show some less-typical convergence curves, including SVM for which the performance gap between methods was extremely large. The curves for non-negative matrix completion show more irregular behavior (including oscillations) because of non-convexity. Finally, we show convergence curves for the total variation minimization in Figure 1d. For this problem, adaptive and accelerated methods were competitive. The adaptive method was superior in the low-precision regime, with the accelerated variant winning out in the high-precision regime. This is largely because total variation involves the gradient operator, which has a large condition number. Nesterov-type acceleration tends to be most effective for these types of poorly-conditioned problems, however the advantages over adaptivity are still slim in the examples shown here.

We note that the advantage of adaptivity is most pronounced for problems where f is non-quadratic, i.e., for problems with a logistic data term. In this case, the Hessian of f varies over the problem domain, and the optimal stepsize for FBS varies with it. For such problems, an adaptive scheme is able to effectively match the stepsize to the local structure of the objective, resulting in fast convergence.

7. CONCLUSION

The forward-backward splitting method is a surprisingly simple way to solve a wide range of optimization problems. Even seemingly complex problems involving non-differentiable objectives (total-variation, support vector machine, sparse regression), and complex constraint sets (semi-definite programming and max-norm regularization, etc...) can be reduced to a sequence of extremely simple steps.

Historically, FBS is most commonly used for simple sparse regression problems despite its much wider applicability. In many common domains, more complex splitting methods involving Lagrange multipliers (such as ADMM [50, 51] and its variants [52, 53]) are more commonly used. However, these methods are often more complex, memory intensive, and computationally burdensome than FBS. Furthermore, FBS has a major advantage over other splitting methods – algorithm parameters like stepsizes and stopping conditions are easily automated. For other splitting methods, it is considerably more difficult to guarantee convergence for adaptive methods [54, 55]. For this reason it is fair to say that FBS is under-utilized for complex problems.

APPENDIX A. CONVERGENCE PROOF FOR NON-MONOTONE LINE SEARCH

We now consider the convergence of the backtracking line search discussed in section 4.4.

Theorem 1. *Suppose that FBS is applied to (1) with convex g and differentiable f . Suppose further that $h = f + g$ is proper, lower semi-continuous, and has bounded level sets. If $\{\tau^k\}$ is bounded below by a positive constant and*

$$(45) \quad f(x^{k+1}) - \hat{f}^k < \langle x^{k+1} - x^k, \nabla f(x^k) \rangle + \frac{1}{2\tau^k} \|x^{k+1} - x^k\|^2$$

then $\lim_{k \rightarrow \infty} h(x^k) = h^*$, where h^* denotes the minimum value of h .

Proof. From the optimality condition (4), we have $0 \in \tau^k \partial g(x^{k+1}) + x^{k+1} - \bar{x}^{k+1}$ and so

$$0 = \tau^k G^{k+1} + x^{k+1} - \bar{x}^{k+1} = \tau^k G^{k+1} + x^{k+1} - (x^k - \tau F^k)$$

for some $G^{k+1} \in \partial g(x^{k+1})$ and $F^k = \nabla f(x^k)$. From this we arrive at

$$(46) \quad x^k - x^{k+1} = \tau^k (G^{k+1} + F^k)$$

and also

$$(47) \quad \langle x^{k+1} - x^k, F^k + G^{k+1} \rangle = -\frac{1}{\tau^k} \|x^{k+1} - x^k\|^2.$$

Now, because g is convex

$$(48) \quad g(x^k) \geq g(x^{k+1}) + \langle x^k - x^{k+1}, G^{k+1} \rangle.$$

Subtracting (48) from (45) and applying (47) yields

$$\begin{aligned} h(x^{k+1}) &= f(x^{k+1}) + g(x^{k+1}) \\ &\leq \hat{f}^k + g(x^k) + \langle x^{k+1} - x^k, F^k + G^{k+1} \rangle \\ &\quad + \frac{1}{2\tau^k} \|x^{k+1} - x^k\|^2 \\ &= \hat{f}^k + g(x^k) - \frac{1}{2\tau^k} \|x^{k+1} - x^k\|^2 \\ (49) \quad &= \hat{h}^k - \frac{1}{2\tau^k} \|x^{k+1} - x^k\|^2. \end{aligned}$$

where $\hat{h}^k = \max\{h^{k-1}, h^{k-2}, \dots, h^{k-\min\{M, k\}}\}$. Note that $\{\hat{h}^k\}$ is a monotonically decreasing bounded sequence, and thus has a limit \hat{h}^* .

We aim to show that \hat{h}^* is the minimal value of h . Observe that $\hat{h}^k = h^{k'}$ for some k' with $k - M \leq k' \leq k$. It is clear from (49) that there must exist a sub-sequence $\{x^{k(i)}\}$ with $h(x^{k(i)}) = \hat{h}^k$ such that

$$(50) \quad \lim_{i \rightarrow \infty} \frac{1}{2\tau^k} \|x^{k(i)+1} - x^{k(i)}\|^2 = 0$$

(otherwise, equation (49) would imply $\hat{h}^k \rightarrow -\infty$ for large k). Note that the level sets of h are assumed to be bounded, and by equation (49) $\{h(x^k)\}$ is bounded as well. By compactness, we may assume without loss of generality that $\{x^{k(i)}\}$ is a convergent sub-sequence of iterates with limit point x^* .

Now, from (46), we have

$$(51) \quad \frac{1}{2\tau^k} \|x^{k(i)+1} - x^{k(i)}\|^2 = \frac{\tau^k}{2} \|G^{k(i)+1} + F^{k(i)}\|^2.$$

Equation (51), together with (50) and the fact that τ^k is bounded away from zero, implies that

$$\lim_{i \rightarrow \infty} \|G^{k(i)+1} + F^{k(i)}\|^2 = 0.$$

Because ∇f is Lipschitz continuous and $\|x^{k(i)+1} - x^{k(i)}\| \rightarrow 0$, we also conclude that $x^{k(i)+1} \rightarrow x^*$ and

$$(52) \quad \lim_{i \rightarrow \infty} \|G^{k(i)+1} + F^{k(i+1)}\|^2 = 0.$$

Note that $G^{k(i)+1} + F^{k(i+1)} \in \partial h(x^{k(i)+1})$. Because the sub-differential of a convex function is continuous¹, (52) implies that $0 \in \partial h(x^*)$, and so x^* is a minimizer of h .

We have shown that $\lim_{k \rightarrow \infty} \hat{h}^k = h(x^*) = h^*$. Because $h^* \leq h(x^k) \leq \hat{h}^k$, we arrive at the conclusion

$$\lim_{k \rightarrow \infty} h(x^k) = h^*.$$

□

REFERENCES

- [1] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach, “Proximal methods for sparse hierarchical dictionary learning,” *Proc. ICML*, 2010.
- [2] J. Duchi and Y. Singer, “Efficient Online and Batch Learning Using Forward Backward Splitting,” *Journal of Machine Learning Research*, vol. 10, pp. 2899–2934, Dec. 2009.
- [3] K. Gregor and Y. LeCun, “Learning fast approximations of sparse coding,” in *Proc. NIPS*, pp. 399–406, 2010.
- [4] A. Szlam, K. Gregor, and Y. LeCun, “Fast Approximations to Structured Sparse Coding and Applications to Object Classification,” in *Lecture Notes in Computer Science*, vol. 7576, pp. 200–213, Springer, 2012.
- [5] M. Hein and S. Setzer, “Beyond spectral clustering-tight relaxations of balanced graph cuts,” in *Proc. NIPS*, pp. 2366–2374, 2011.
- [6] P. Olsen, F. Oztoprak, J. Nocedal, and S. Rennie, “Newton-like methods for sparse inverse covariance estimation,” in *Proc. NIPS*, pp. 764–772, 2012.
- [7] H. Jégou, T. Furon, and J.-J. Fuchs, “Anti-sparse coding for approximate nearest neighbor search,” *arXiv:1110.3767v2*, Oct. 2011.
- [8] M. Figueiredo, R. Nowak, and S. Wright, “Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 586–597, 2007.
- [9] P. Combettes and J. Pesquet, “Proximal splitting methods in signal processing,” in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer Optimization and Its Applications, pp. 185–212, Springer, 2011.
- [10] P. L. Combettes and V. R. Wajs, “Signal recovery by proximal forward-backward splitting,” *Multiscale Modeling and Simulation*, vol. 4, no. 4, pp. 1168–1200, 2005.
- [11] D. Bertsekas, “On the goldstein-levitin-polyak gradient projection method,” *Automatic Control, IEEE Transactions on*, vol. 21, no. 2, pp. 174–184, 1976.
- [12] A. A. Goldstein, “Convex programming in hilbert space,” *Bulletin of the American Mathematical Society*, vol. 70, pp. 709–710, 09 1964.
- [13] E. S. Levitin and B. T. Polyak, “Constrained minimization methods,” *USSR Computational mathematics and mathematical physics*, vol. 6, no. 5, pp. 1–50, 1966.
- [14] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society, Series B*, vol. 58, pp. 267–288, 1994.
- [15] J. Duchi, S. S. Schwartz, Y. Singer, and T. Chandra, “Efficient projections onto the ℓ_1 -ball for learning in high dimensions,” in *Proc. of the 25th international conference on Machine learning, ICML '08*, (New York, NY, USA), pp. 272–279, ACM, 2008.

¹More formally, the convex sub-differential is a multi-valued function which is upper semicontinuous in a topological sense. See [56, 57].

- [16] E. van den Berg and M. P. Friedlander, “SPGL1: A solver for large-scale sparse reconstruction,” June 2007. <http://www.cs.ubc.ca/labs/scl/spgl1>.
- [17] E. van den Berg and M. P. Friedlander, “Probing the pareto frontier for basis pursuit solutions,” *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.
- [18] S. Cotter, B. Rao, K. Engan, and K. Kreutz-Delgado, “Sparse solutions to linear inverse problems with multiple measurement vectors,” *Signal Processing, IEEE Transactions on*, vol. 53, pp. 2477–2488, July 2005.
- [19] C. Studer, T. Goldstein, W. Yin, and R. Baraniuk, “Democratic representations,” *arXiv:1401.3420*, Jan. 2014.
- [20] R. Keshavan, A. Montanari, and S. Oh, “Low-rank matrix completion with noisy observations: A quantitative comparison,” in *The 47th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1216–1222, 2009.
- [21] M. Davenport, Y. Plan, E. van den Berg, and M. Wootters, “1-Bit matrix completion,” *arXiv:1209.3672*, Sep. 2013.
- [22] L. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica. D.*, vol. 60, pp. 259–268, 1992.
- [23] A. Chambolle, “An algorithm for total variation minimization and applications,” *J. Math. Imaging Vis.*, vol. 20, no. 1-2, pp. 89–97, 2004.
- [24] A. Beck and M. Teboulle, “Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems,” *Image Processing, IEEE Transactions on*, vol. 18, pp. 2419–2434, Nov 2009.
- [25] A. Smola and B. Schölkopf, “A tutorial on support vector regression,” *Statistics and Computing*, vol. 14, no. 3, pp. 199–222, 2004.
- [26] M. Goemans and D. Williamson, “Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming,” *J. ACM*, vol. 42, no. 6, pp. 1115–1145, 1995.
- [27] E. Candes, T. Strohmer, and V. Voroninski, “PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming,” *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [28] D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, pp. 788–791, Oct. 1999.
- [29] A. N. Langville, C. D. Meyer, R. Albright, J. Cox, and D. Duling, “Algorithms, initializations, and convergence for the nonnegative matrix factorization,” *CoRR*, vol. abs/1407.7299, 2014.
- [30] D. D. Lee and H. S. Seung, “Algorithms for non-negative matrix factorization,” in *In NIPS*, pp. 556–562, MIT Press, 2000.
- [31] J. Lee, B. Recht, R. Salakhutdinov, N. Srebro, and J. Tropp, “Practical large-scale optimization for max-norm regularization,” in *Proc. NIPS*, 2010.
- [32] J. Barzilai and J. M. Borwein, “Two-point step size gradient methods,” *IMA J Numer Anal*, vol. 8, pp. 141–148, January 1988.
- [33] E. G. Birgin, J. . M. Martinez, and M. Raydan, “Nonmonotone spectral projected gradient methods on convex sets,” *SIAM Journal on Optimization*, pp. 1196–1211, 2000.
- [34] Z. Wen, W. Yin, and D. Goldfarb, “On the convergence of an active set method for ℓ_1 minimization,” tech. rep., Rice University, 2010.
- [35] T. Goldstein and S. Setzer, “High-order methods for basis pursuit,” *Under review (available as UCLA CAM Report)*, 2010.
- [36] S. Wright, R. Nowak, and M. Figueiredo, “Sparse reconstruction by separable approximation,” *Signal Processing, IEEE Transactions on*, vol. 57, pp. 2479–2493, July 2009.
- [37] B. Zhou, L. Gao, and Y.-H. Dai, “Gradient methods with adaptive step-sizes,” *Comput. Optim. Appl.*, vol. 35, pp. 69–86, Sept. 2006.
- [38] R. Fletcher, “On the Barzilai-Borwein Method,” in *Optimization and Control with Applications*, pp. 235–256, Springer, 2005.
- [39] D. Bertsekas, “Projected newton methods for optimization problems with simple constraints,” *SIAM Journal on Control and Optimization*, vol. 20, no. 2, pp. 221–246, 1982.
- [40] J. Lee, Y. Sun, and M. Saunders, “Proximal newton-type methods for convex optimization,” in *Advances in Neural Information Processing Systems 25* (P. Bartlett, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, eds.), pp. 836–844, 2012.

- [41] J. Bioucas-Dias and M. Figueiredo, "A new TwIST: Two-step iterative shrinkage/thresholding algorithms for image restoration," *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 2992–3004, Dec. 2007.
- [42] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal of Imaging Sciences*, vol. 2, pp. 183–202, 2009.
- [43] B. O'Donoghue and E. Candès, "Adaptive restart for accelerated gradient schemes." <http://www-stat.stanford.edu/~candes/publications.html>, 2012. Manuscript.
- [44] R. Burachik, L. Mauricio, G. Drummond, A. Iusem, and E. Castorina, "Full convergence of the steepest descent method with inexact line searches," *Optimization*, vol. 32, pp. 137–146, 1995.
- [45] K. Kiwiel and K. Murty, "Convergence of the steepest descent method for minimizing quasi convex functions," *Journal of Optimization Theory and Applications*, vol. 89, pp. 221–226, 1996.
- [46] L. Grippo, F. Lampariello, and S. Lucidi, "A Nonmonotone Line Search Technique for Newton's Method," *SIAM Journal on Numerical Analysis*, vol. 23, no. 4, pp. 707–716, 1986.
- [47] Z. Hongchao and W. Hager, "A nonmonotone line search technique and its application to unconstrained optimization," *SIAM J. Optim.*, vol. 14, pp. 1043–1056, 2004.
- [48] S. Becker, J. Bobin, and E. Candès, "Nesta: A fast and accurate first-order method for sparse recovery," *Technical Report of Stanford University*, Apr 2009.
- [49] E. Hale, W. Yin, and Y. Zhang, "A fixed-point continuation method for ℓ_1 -regularized minimization with applications to compressed sensing," *CAAM Technical Report*, vol. TR07, 2007.
- [50] R. Glowinski and P. L. Tallec, *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1989.
- [51] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends in Machine Learning*, 2010.
- [52] E. Esser, X. Zhang, and T. Chan, "A general framework for a class of first order primal-dual algorithms for TV minimization," *UCLA CAM Report 09-67*, 2009.
- [53] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Convergence*, vol. 40, no. 1, pp. 1–49, 2010.
- [54] B. He, H. Yang, and S. Wang, "Alternating direction method with self-adaptive penalty parameters for monotone variational inequalities," *Journal of Optimization Theory and Applications*, vol. 106, no. 2, pp. 337–356, 2000.
- [55] T. Goldstein, E. Esser, and R. Baraniuk, "Adaptive primal-dual hybrid gradient methods for saddle-point problems," *Available at Arxiv.org (arXiv:1305.0546)*, 2013.
- [56] P. Chakrabarty, A.K. and Shunmugaraj and C. Zalinescu, "Continuity properties for the sub-differential and ϵ -subdifferential of a convex function and its conjugate," *Journal of Convex Analysis*, vol. 14, pp. 479–514, 2007.
- [57] H. Bauschke and P. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2011.